

Intel® VMD NVMe RAID)

Quick Configuration Guide

A quickstart guide detailing configuration rules and procedures

Rev 1.2

March 2020

<Blank page>

Document Revision History

Date Published	Revision	Changes
December 2019	1.0	Initial release
December 2019	1.1	Corrections made to the list of supported OS and VMD definition. Section 2 ¹ Preparing the server hardware and drivers Appendix B - Added.
March 2020	1.2	Removed Secure Boot restriction.

Disclaimers

Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 800-548-4725 or by visiting www.intel.com/design/literature.htm

Intel, the Intel logo, Intel® Virtual RAID on CPU, Intel® Rapid Storage Technology, enterprise are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

© 2018 -2019 Intel Corporation

Warnings

Heed safety instructions: Before working with your server product, whether you are using this guide or any other resource as a reference, pay close attention to the safety instructions. You must adhere to the assembly instructions in this guide to ensure and maintain compliance with existing product certifications and approvals. Use only the described, regulated components specified in this guide. Use of other products/components will void the UL listing and other regulatory approvals of the product and most likely result in noncompliance with product regulations in the region(s) in which the product is sold.

System power on/off: The power button DOES NOT turn off the system AC power. To remove all power from the system, you must unplug the AC power cord from the wall outlet. Make sure the AC power cord is unplugged before you open the chassis, add, or remove any components.

Hazardous conditions, devices and cables: Hazardous electrical conditions may be present on power, telephone, and communication cables. Turn off the server and disconnect the power cord, telecommunications systems, networks, and modems attached to the server before opening it. Otherwise, personal injury or equipment damage can result.

Installing or removing jumpers: A jumper is a small plastic-encased conductor that slips over two jumper pins. Some jumpers have a small tab on top that you can grip with your fingertips or with a pair of fine needle-nosed pliers. If your jumpers do not have such a tab, take care when using needle-nosed pliers to remove or install a jumper; grip the narrow sides of the jumper with the pliers, never the wide sides. Gripping the wide sides can damage the contacts inside the jumper, causing intermittent problems with the function controlled by that jumper. Take care to grip with, but not squeeze, the pliers or other tool you use to remove a jumper, or you may bend or break the pins on the board.

Electrostatic discharge (ESD) and ESD protection: ESD can damage disk drives, boards, and other parts. We recommend that you perform all procedures in this chapter only at an ESD workstation. If one is not available, provide some ESD protection by wearing an anti-static wrist strap attached to chassis ground, any unpainted metal surface on your server when handling parts.

ESD and handling boards: Always handle boards carefully. They can be extremely sensitive to ESD. Hold boards only by their edges. After removing a board from its protective wrapper or from the server, place the board component side up on a grounded, static-free surface. Use a conductive foam pad if available but not the board wrapper. Do not slide board over any surface.

Preface

Introduction

This guide provides concise instructions for creating RAID configurations using VROC HII Configuration Utility. The utility runs prior to loading the operating system (OS) and can be used to prepare a RAID volume before an OS is loaded onto it.

Guide Organization

This guide includes the following:

Section 1 - Product Overview

Provides the product overview, supported hardware and operating systems.

Section 2 - Preparing the server hardware and drivers

Provides the steps required to prepare the system to support Intel® VROC RAID configurations

Section 3 - Intel® Volume Management Devices (Intel® VMD)

Provides definition of the Intel® VMD feature and instructions to enable it.

Section 4 - Creating Intel® VROC (VMD NVMe RAID) volumes

Step by step guide and screenshots to create Intel® VROC RAID configurations.

Table of Contents

Introduction.....	i
Guide Organization.....	i
1. Overview.....	4
1.1 SupportedHardware.....	4
1.2 SupportedOperatingSystems*.....	4
2. Preparing Server Hardware and Drivers.....	5
2.1 Server System.....	5
2.2 Intel VROC Hardware Key.....	5
2.3 NVMe* drives.....	5
2.4 PCIe switches and retimers.....	5
2.5 Drivers.....	6
3. Intel® Volume Management Device (Intel® VMD).....	8
3.1 Enabling the Intel® VMD ports.....	8
4. Creating Intel® VROC (VMD NVMe RAID) volumes.....	11
4.1 Accessing the Intel® VROC (VMD NVMe RAID) HII Configuration Utility.....	11
4.2 Creating RAID Volumes Using the Intel® VROC (VMD NVMe RAID) HII Configuration Utility.....	11
Appendix A. Glossary.....	15
Appendix B. Updating Firmware on Intel NVMe drives using the Intel* SSD Data Center Tool.....	16

List of Tables

Table 1. VROC driver requirements.....	7
--	---

List of Figures

Figure 1. RSTe NVMe release note documentation style.....	6
Figure 2. Steps to enter the Volume Management Device Port Window.....	9
Figure 3. VMD port Example A.....	9
Figure 4. VMD port Example B.....	10
Figure 5. Steps to access the Intel® VROC configuration utility.....	11
Figure 6. Intel® VROC Managed VMD menu.....	12
Figure 7. Create RAID Volume menu.....	12
Figure 8. Completed Create RAID Volume menu.....	13
Figure 9. Intel® VROC Managed VMD menu.....	14
Figure 10. List of Intel® NVMe drives shown by the isdct tool.....	16
Figure 11. Updating an NVMe drive.....	16

<Blank page>

1. Overview

Intel® VROC (VMD NVMe RAID) is a powerful, reliable, and affordable hybrid RAID solution that uses the built-in Volume Management Devices (VMD) and Intel® RSTe drivers to create and manage RAID configurations. Intel® VROC is only available when the system is configured for UEFI boot mode. To enable the Intel® VROC hardware key must be inserted into the motherboard, and the appropriate processor Virtual Management Devices must be enabled in BIOS Setup.

Intel® VROC provides added performance and reliability for supported systems equipped with NVMe Express* (NVMe*) drives to enable an optimal enterprise storage solution. Intel® VROC offers data protection by enabling RAID in pre-boot environment that can be used in operating systems like Microsoft Windows*, Red Hat Linux* and SUSE Linux*.

1.1 Supported Hardware

This guide covers server products and solutions based on the following Intel® Server Board product families:

- Intel® Server Board S2600WF product family (a.k.a Wolf Pass)
- Intel® Server Board S2600BP product family (a.k.a Buchanan Pass)
- Intel® Server Board S2600ST product family (a.k.a Sawtooth Pass)

The guide supports RAID configurations created using validated NVMe* drives connected through a supported backplane. To increase the number of supported NVMe* drives, approved PCIe* switches and retimers are also supported.

Note: For a list of supported devices, consult the Tested Hardware List ([THOL](#)).

1.2 Supported Operating Systems*

The Intel® VROC (VMD NVMe RAID) solution supports the following Operating Systems, either within drivers or out of the box drivers:

- Microsoft Windows Server 2012*R2
- Microsoft Windows Server 2016*
- Microsoft Windows Server 2019*
- Red Hat Enterprise Linux 7.3
- Red Hat Enterprise Linux 7.4
- Red Hat Enterprise Linux 7.5
- Red Hat Enterprise Linux 7.6
- Red Hat Enterprise Linux 8.0
- SUSE Linux Enterprise Server 12 SP3
- SUSE Linux Enterprise Server 15
- SUSE Linux Enterprise Server* 15 SP1

Note: Some of these Operating Systems might not be supported on a particular server system. The following list contains the OS tests for each system:

- [Tested Operating Systems for Intel® Server Board S2600WF Family](#)
 - [Tested Operating Systems for Intel® Server Board S2600ST Family](#)
 - [Tested Operating Systems for Intel® Server Board S2600BP Family](#)
-

2. Server System

2.1 Server System

Intel recommends that the server system be installed with the latest firmware package that includes the BIOS, ME, FD, BMC, and FRUSDR utility. The best way to achieve this is to download and install the latest published System Update Package (SUP) corresponding to the motherboard being used, which may be found at the [Intel Download Center](#). In the search box near the top of the page in the motherboard model number (e.g., S2600WFT), then click the magnifying glass icon. Set the "Download Type" dropdown to "BIOS" or "Firmware" (depending on which option is available) and the available SUP versions will be displayed. The specific package contains installation instructions.

2.2 Intel VROC Hardware Key

In order to enable Intel® VROC a Hardware Key must be installed on the Intel® motherboard. There are three separate keys supporting different RAID levels:

- Intel® VROC Standard (MM#951605): Enables RAID levels Q 1 and 10. Different NVMe drive manufacturers are supported. Consult [the HOL](#) for a full list.
- Intel® VROC Premium (MM#951606): RAID levels Q 1, 10 and 5. Different NVMe drive manufacturers are supported. Consult [the HOL](#) for a full list.
- Intel® VROC Intel SSD Only (MM#956822): Enables RAID levels Q 1, 10 and 5 but only the Intel NVMe drives are supported.

Refer to the corresponding system BIOS document of your Intel® motherboard for VROC Hardware Key installation instructions.

2.3 NVMe* drives

Intel® VROC only supports the U.2 form factor NVMe drives connected through supported backplanes. Intel recommends that the drives are of the latest firmware version, and makes available the [Intel* SSD Data Center Tool](#) utility in order to update the firmware on Intel SSD DC family drives. Other vendors provide similar tools.

The Intel® SSD Data Center Tool runs on top of Windows* or Linux operating systems and can be used when the drives are behind the VMD controllers as well as after a RAID configuration has been created. For more detailed information on how to update the firmware on the Intel NVMe drives, see the [Intel* SSD Data Center Tool](#) refer to Appendix B.

2.4 PCIe switches and retimers

Intel recommends having the latest firmware installed on the PCIe switches, ensuring that the BMC is aware of such cards. If the SUP (mentioned in Section 2.1) was applied after the cards were installed, then the BMC is already aware. In other cases, or if the cards were moved to a different PCIe slot, the BMC needs to be made aware of such changes by booting the system into the EFI shell and running the `sdr_update_noprompt.nsh` script. This script is part of the SUP.

The `sdr_update_noprompt.nsh` script handles the drive LED management. In other words, who handles the create, fault and rebuild LEDs. LED management is handled in a different way when VMD is enabled than when VMD disabled.

LED management is set to **VMD disabled** by default but since VMD must be enabled in order to have Intel® VROC, it must be set to **VMD enabled** the first time the PCIe switch is installed. This setting is persistent, so it only needs to be set once. Intel provides an LED Management toggle tool to change this setting.

The [Firmware Package and LED Mode Toggle Tool for Intel® PCIe* Switches](#) includes both, the firmware package and the LED Management toggle tool. The package has readme files with instructions for each component.

Note: Not configuring LED management correctly will make the locate and fault drive LEDs not to work and make difficult to identify and replace a faulty drive.

2.5 Drivers.

The Intel® VROC solution has two driver components, the pre-boot, or UEFI driver, and the OS driver. The UEFI driver is embedded in the system BIOS and referred to as the **RSTe NVMe UEFI driver**, while the OS driver must be loaded at the OS installation time and referred to as the **RSTe OS driver**; both contain version numbers. The OS driver version can be equal or later than the UEFI driver version but vice versa excepting for the Linux OS. If the system BIOS version is up to date while the prior BIOS version has a newer RSTe NVMe UEFI driver embedded, the RSTe OS driver must be updated first.

Note: On a system with the Windows OS installed and an Intel® VROC virtual drive present, the **Windows* RSTe OS driver** version must not be prior to the **RSTe NVMe UEFI driver** version, otherwise unexpected behavior may happen.

The RSTe OS driver can be found on the [Intel Download Center](#). Enter vroc into the field near the magnifier icon, and it will show both Windows and Linux driver links.

The RSTe NVMe UEFI driver version embedded in the BIOS is documented in the corresponding BIOS Release Notes for the SUP Package, look at the BIOS COMPONENTS/CONTENTS

Below is an example of how the RSTe NVMe driver version is documented in the Release Notes file.

```

Release Notes_R00.01.0014.txt - Notepad
File Edit Format View Help
S2600BP Family
S2600ST Family

=====
BIOS COMPONENTS/CONTENTS
=====
Processors supported:
Intel(R) Xeon Scalable Family Processor

Microcode versions:
CPUID      Version      Status
0x50653    0x01000144   (Xeon Scalable B1)
0x50654    0x0200004D   (Xeon Scalable H0)

SATAAHCI:
v2.00i
RSTeSATAEfi:
v5.4.0.1039
RSTeSATALegacy:
v5.4.0.1039
RSTesSATAEfi:
v5.4.0.1039
RSTesSATALegacy:
v5.4.0.1039
RSTe NVMe UEFI drivers:
v5.4.0.1039
ESRT2 UEFI drivers:
05120000
    
```

Figure 1. RSTe NVMe release note documentation style

For the UEFI driver version in the above example, the 5.4 or later RSTe OS driver must be used, but nothing earlier than 5.4 would be supported.

Intel® VROC Quick Configuration Guide

In the first iteration of this guide, the latest BIOS version available was 02.01.0009. Refer to the following table to choose the best OS driver to use, depending on the system BIOS version used.

Table 1. VROC driver requirements

System BIOS version	VROC NVMe UEFI driver version	VROC Windows driver version required	VROC Linux driver version required
00.01.0004	5.2.1.1001	6.0.0.1357 (or later)	5.4 (or later)
00.01.0009	5.2.1.1001	6.0.0.1357 (or later)	5.4 (or later)
00.01.0012	5.3.0.1052	6.0.0.1357 (or later)	5.4 (or later)
00.01.0013	5.3.0.1052	6.0.0.1357 (or later)	5.4 (or later)
00.01.0014	5.4.0.1039	6.0.0.1357 (or later)	5.4 (or later)
00.01.0015	5.4.0.1039	6.0.0.1357 (or later)	5.4 (or later)
00.01.0016	5.5.0.1028	6.0.0.1357 (or later)	6.0 2019.04.23 (or later)
02.01.0008	6.0.0.1024	6.0.0.1357 (or later)	6.0 2019.04.23 (or later)
02.01.0009	6.1.0.1017	6.1.0.1247 (or later)	6.0 2019.04.23 (or later)

3. Enabling and Configuring Intel® VMD Ports

Intel® VMD is a new feature introduced with the Intel® Xeon® processor Scalable family. This native feature supported backplane. This section describes how to enable and configure this functionality.

Each member of the Intel® Xeon® processor Scalable family has three Intel® VMD domains (numbered one, two, and three). Each Intel® VMD domain manages 16 PCIe* lanes divided into four Intel® VMD ports (named A, B, C, and D) consisting of four PCIe* lanes each. Some of the Intel® VMD ports are routed to specific risers and slots in the system, while other Intel® VMD ports are routed to specific chipset uplinks, SAS modules, or onboard Oculink connectors. This routing is fixed (non-configurable) please refer to the relevant motherboard's Technical Product Specification for details on this routing.

Note: Intel® VMD ports routed to specific chipset uplinks cannot be used for Intel® VROC (VMD NVMe RAID) configurations

There are two different ways to connect NVMe* drives to Intel® VMD ports:

1. NVMe* drives connect directly to VMD ports in a 1:1 fashion using the onboard Oculink* ports or by using PCIe* retimers. Each NVMe* drive then uses four dedicated PCIe* lanes, providing full bandwidth to each drive
2. Several NVMe* drives connect to one VMD port by using either 8x4 or 8x8 PCIe* switches. Eight PCIe* lanes are then shared by the NVMe* ports providing limited bandwidth to each drive

The RAID volume may be used as a bootable drive only if all drives in the RAID volume are connected to a single Intel® VMD domain. It is possible to create RAID volumes spanning multiple Intel® VMD domains, however such RAID configurations cannot be made bootable.

3.1 Enabling the Intel® VMD ports

By default, all Intel® VMD ports are disabled in BIOS Setup. For Intel® VROC (VMD NVMe RAID), the appropriate Intel® VMD ports must be enabled by selecting the following menu options in order after entering BIOS Setup

1. Advanced
2. PCI Configuration
3. Volume Management Device

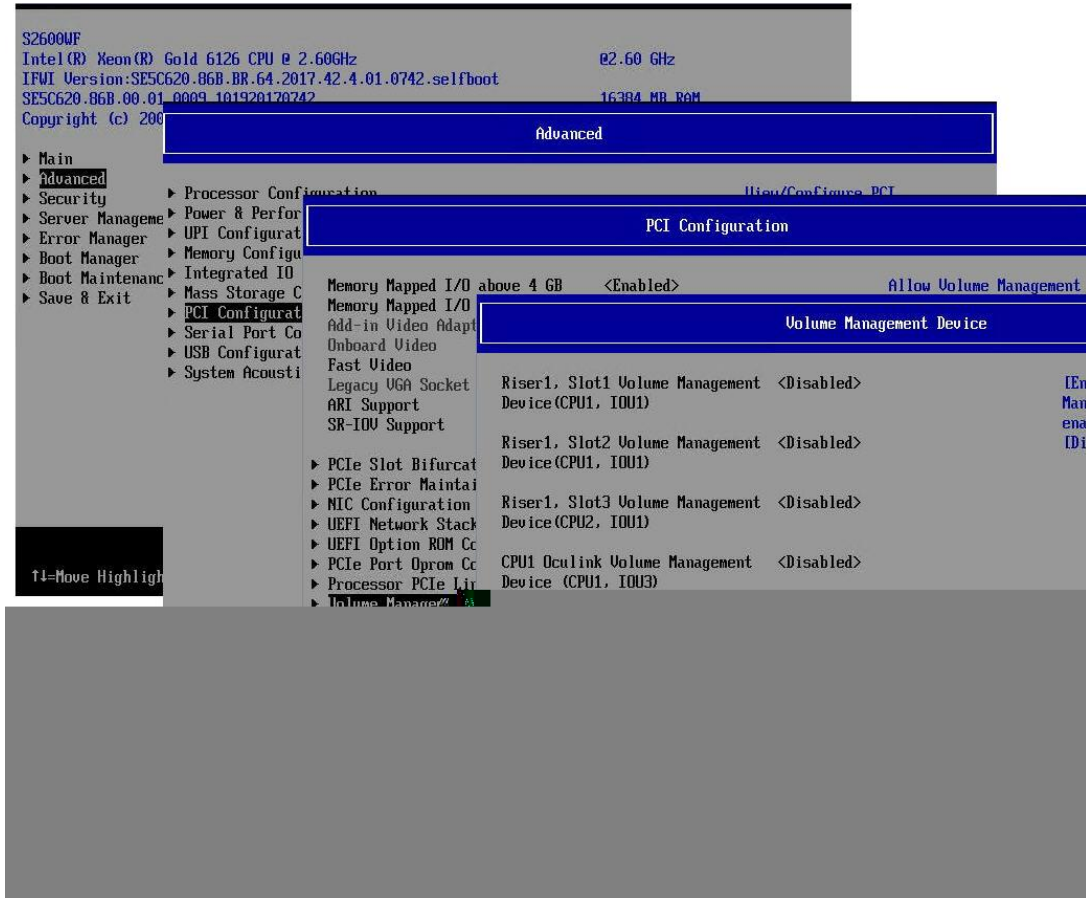


Figure 2. Steps to enter the Volume Management Device Port Window

Example A: If an R2000WF system is being used and four NVMe* drives are being connected to the four onboard Oculink connectors, VMD ports 3A, 3B, 3C and 3D must be enabled.

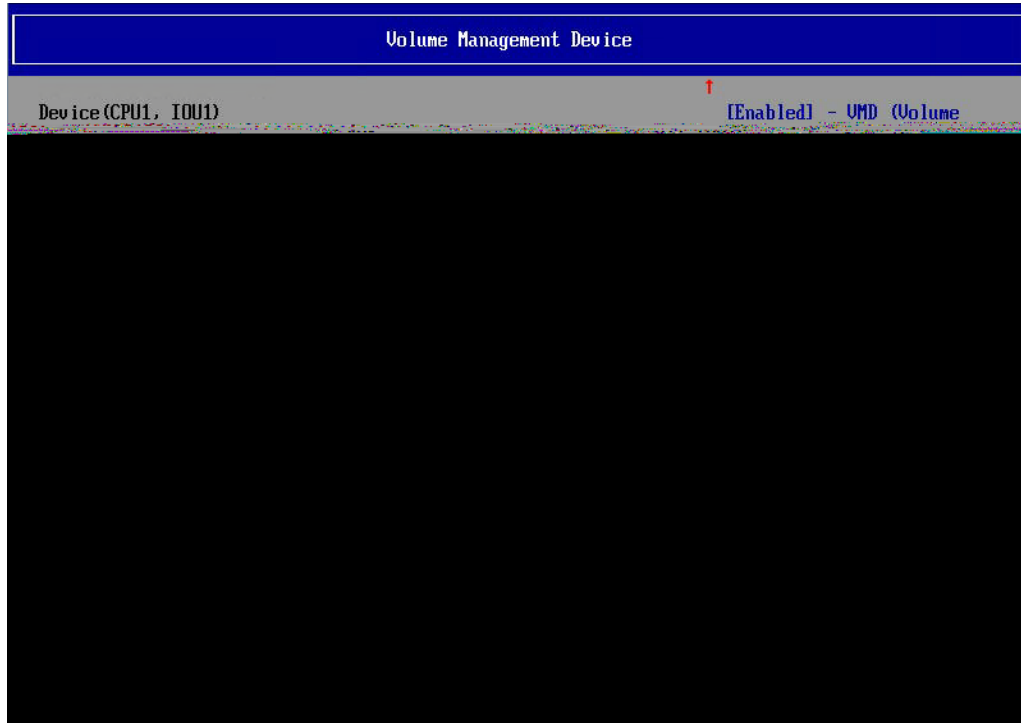


Figure 3. VMD port Example A

Intel® VROC Quick Configuration Guide

1. After selecting Intel® Virtual RAID on CPU, the Intel® VROC (VMD NVMe RAID) menu screen appears, showing the installed Upgrade Key and the configured Intel® VROC RAID volumes (if any).
2. Select All Intel® VMD Controllers. The Intel® VROC managed VMD menu will appear showing all Intel® VMD controllers with the corresponding NVMe* drives.

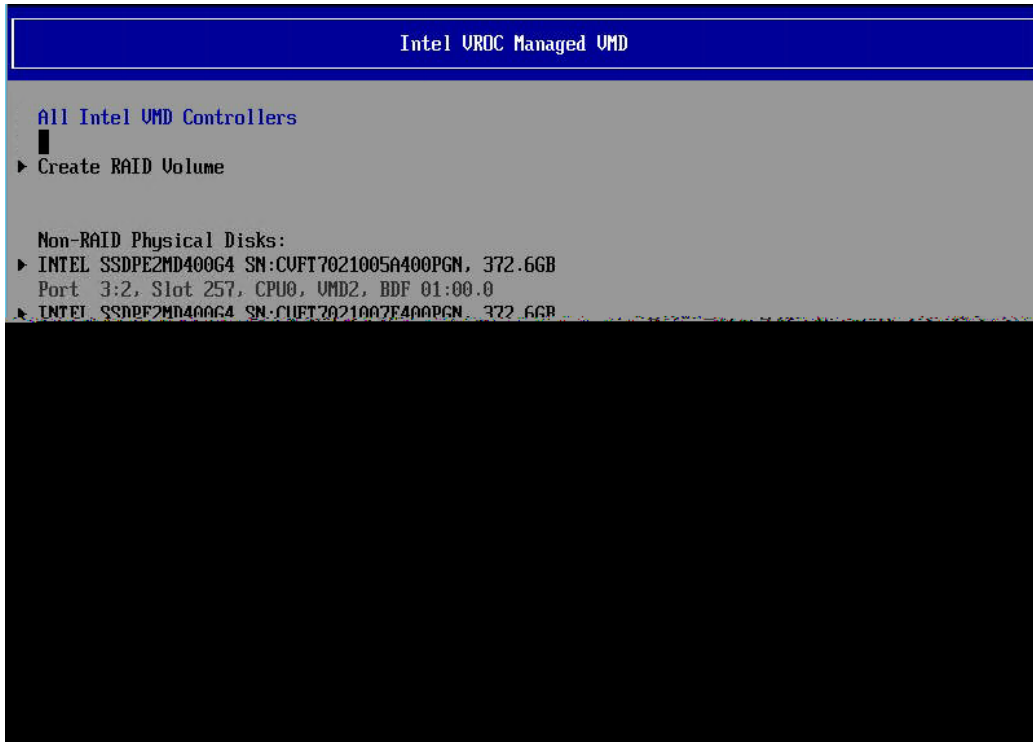


Figure 6. Intel® VROC Managed VMD menu

3. Select Create RAID Volume. The Create RAID Volume menu will appear.

Figure 7. Create RAID Volume menu

4. Fill in the volume name (optional), then select the desired RAID level depending on the number of

Appendix A. Glossary

Term	Definition
Intel® RSTe	Intel®Rapid StorageTechnology enterprise
NVMe*	NVM Express*
Intel® VMD	Intel®Volume Management Devices
Intel® VROC	Intel® Virtual RAID on CPU

Appendix B. Updating Firmware on Intel NVMe drives using the Intel* SSD Data Center Tool

- Download and install the [Intel* SSD Data Center Tool](#)
For Windows* OS, the binaries are installed within the `C:\isdct` folder.
For Linux* OS, the binaries are installed within the `/usr/bin` directory while other support files are installed within the `/etc` and `/usr/lib` directories.
- Run the next command `isdct show -intelssd`
All the installed NVMe drives will be listed along with information like serial number, model number, firmware version, index number, etc (see Figure 10).
Showing the available version.

```
- Intel SSD DC P4501 Series PHLF7504007M2P0LGN -
Bootloader : 0136
DevicePath : \\.\PHYSICALDRIVE5
DeviceStatus : Healthy
Firmware : QDV101D1
FirmwareUpdateAvailable : The selected Intel SSD contains current firmware as of this tool release.
Index : 3
ModelNumber : INTEL SSDPE7KX020T7
ProductFamily : Intel SSD DC P4501 Series
SerialNumber : PHLF7504007M2P0LGN

- Intel SSD DC S3500 Series BTWM614205AU120B -
Bootloader : Property not found
DevicePath : CSMI_C0R0P1P1
DeviceStatus : Healthy
Firmware : G2010140
FirmwareUpdateAvailable : G2010150
Index : 4
ModelNumber : INTEL SSDSCKHB120G4
ProductFamily : Intel SSD DC S3500 Series
SerialNumber : BTWM614205AU120B
```

Figure 10. List of Intel® NVMe drives shown by the isdct tool

- Take note of the Index number for all the drives requiring the firmware update. In the example shown above, the drive containing index number 4 requires the firmware update.
- Run the next command `isdct load -intelssd <index number>`
Confirm when prompted
Wait for the firmware update (see Figure 11).

```
C:\isdct>isdct load -intelssd 4
WARNING! You have selected to update the drives firmware!
Proceed with the update? (Y|N): y
Updating firmware...

- Intel SSD DC S3500 Series BTWM614205AU120B -

Status : Firmware Updated Successfully. Please reboot the system.
```

Figure 11. Updating an NVMe drive.

